On Improved Bounds for Probability Metrics and f-Divergences

Igal Sason

Abstract

Derivation of tight bounds for probability metrics and f-divergences is of interest in information theory and statistics. This paper provides elementary proofs that lead, in some cases, to significant improvements over existing bounds; they also lead to the derivation of some existing bounds in a simplified way. The inequalities derived in this paper relate between the Bhattacharyya parameter, capacitory discrimination, chi-squared divergence, Chernoff information, Hellinger distance, relative entropy, and the total variation distance. The presentation is aimed to be self-contained.

arXiv:1403.7164v2 [cs.IT] 31 Mar 2014

I. Introduction

Derivation of tight bounds for probability metrics and f-divergences is of interest in information theory and statistics, as is reflected from the bibliography of this paper and references therein. Following previous work in this area, elementary proofs are used in this paper for the derivation of bounds. In some cases, existing bounds are re-derived in a simplified way, and in some others, significant improvements over existing bounds are obtained.

The paper is structured as follows: the bounds and their proofs are introduced in Section II, followed by various discussions and remarks that link the new bounds to the literature. This section is separated into four parts: the first part introduces bounds on the Hellinger distance and Bhattacharyya parameter in terms of the total variation distance and the relative entropy (see Section II-A), the second part introduces a lower bound on the Chernoff information in terms of the total variation distance (see Section II-B), the third part provides bounds on the chi-squared divergence and some related inequalities on the relative entropy and total variation distance (see Section II-C), and the last part considers bounds on the capacitory discrimination (see Section II-D). A summary, which outlines the contributions made in this work, is provided in Section III.

Preliminaries

We introduce, in the following, some preliminary material that is essential to make the presentation self-contained.

Definition 1: Let f be a convex function defined on $(0,\infty)$ with $f(1)=0$, and let P and Q be two probability distributions defined on a common set $\mathcal{X}$. The *f-divergence* of P from Q is defined by

$$D_f(P||Q) \triangleq \sum_{x\in\mathcal{X}} Q(x)\, f\left(\frac{P(x)}{Q(x)}\right) \tag{1}$$

where sums may be replaced by integrals. Here we take

$$0f\Big(\frac{0}{0}\Big)=0, \quad f(0)=\lim_{t\to 0^+} f(t), \quad 0f\Big(\frac{a}{0}\Big)=\lim_{t\to 0^+} tf\Big(\frac{a}{t}\Big)=a\lim_{u\to\infty}\frac{f(u)}{u},\ \forall\, a>0.$$

Definition 2: An f-divergence is said to be *symmetric* if the equality $f(x)=xf\big(\frac{1}{x}\big)$ holds for every $x>0$. This requirement on f implies that $D_f(P||Q)=D_f(Q||P)$ for every pair of probability distributions P and Q.

From [13] and [15, Corollary 5.4], the following lower bound holds for a symmetric f-divergence:

$$D_f(P||Q) \geq \big(1-d_{\mathrm{TV}}(P,Q)\big)\, f\left(\frac{1+d_{\mathrm{TV}}(P,Q)}{1-d_{\mathrm{TV}}(P,Q)}\right). \tag{2}$$

Definition 3: Let P and Q be two probability distributions defined on a set $\mathcal{X}$. The *total variation distance* between P and Q is defined by

$$d_{\mathrm{TV}}(P,Q) \triangleq \sup_{\text{Borel } A\subseteq\mathcal{X}} |P(A)-Q(A)| \tag{3}$$

where the supremum is taken over all the Borel subsets A of $\mathcal{X}$.

I. Sason is with the Department of Electrical Engineering, Technion–Israel Institute of Technology, Haifa 32000, Israel (e-mail: sason@ee.technion.ac.il).

If $\mathcal{X}$ is a countable set, (3) is simplified to

$$d_{\mathrm{TV}}(P,Q) = \frac{1}{2}\sum_{x\in\mathcal{X}} |P(x)-Q(x)| = \frac{||P-Q||_1}{2} \tag{4}$$

and in the continuous case, probability mass functions are replaced by probability density functions, and sums are replaced by integrals. The total variation distance is a symmetric f-divergence where $f(t) = \frac{1}{2}|t-1|$ for $t \in \mathbb{R}$.

Definition 4: Let P and Q be two probability distributions that are defined on a common set $\mathcal{X}$. The *Hellinger distance* and the *Bhattacharyya parameter* between P and Q are, respectively, given by

$$d_{\mathrm{H}}(P,Q) \triangleq \left(\frac{1}{2}\sum_{x\in\mathcal{X}}\Big(\sqrt{P(x)}-\sqrt{Q(x)}\Big)^2\right)^{\frac{1}{2}} \tag{5}$$

$$Z(P,Q) \triangleq \sum_{x\in\mathcal{X}} \sqrt{P(x)\,Q(x)}\,. \tag{6}$$

The three measures in (3)–(6) are bounded between 0 and 1. Also, it is easy to verify that

$$d_{\mathrm{H}}(P,Q) = \sqrt{1-Z(P,Q)}. \tag{7}$$

The square of the Hellinger distance is a symmetric f-divergence since the convex function

$$f(x) = \frac{1}{2}\,(1-\sqrt{x})^2, \quad x \geq 0 \tag{8}$$

satisfies the equality $f(x) = x f\big(\frac{1}{x}\big)$ for every $x > 0$ with $f(1) = 0$, and from (1) and (5)

$$\big(d_{\mathrm{H}}(P,Q)\big)^2 = D_f(P||Q). \tag{9}$$

Definition 5: The *Chernoff information* and *relative entropy* (a.k.a. information divergence or Kullback-Leibler distance) between two probability distributions P and Q defined on a common set $\mathcal{X}$ are, respectively, given by

$$C(P,Q) \triangleq -\min_{\theta\in[0,1]} \log\left(\sum_{x\in\mathcal{X}} P(x)^\theta\, Q(x)^{1-\theta}\right) \tag{10}$$

$$D(P||Q) \triangleq \sum_{x\in\mathcal{X}} P(x)\log\left(\frac{P(x)}{Q(x)}\right) \tag{11}$$

where throughout this paper, the logarithms are on base e.

Note that, in general, $C(P,Q), D(P||Q) \in [0,\infty]$, $C(P,Q) = C(Q,P)$, and $D(P||Q) \neq D(Q||P)$. The relative entropy is an asymmetric f-divergence where $f(t) = t\log(t)$ for $t > 0$ is a convex function with $f(1) = 0$.

Proposition 1: For two probability distributions P and Q that are defined on a common set $\mathcal{X}$

$$d_{\mathrm{TV}}(P,Q) \leq \sqrt{2}\,d_{\mathrm{H}}(P,Q) \leq \sqrt{D(P||Q)}. \tag{12}$$

The left-hand side of (12) is proved in [23, p. 99], and the right-hand side is proved in [23, p. 328].

The Chernoff information, $C(P,Q)$, is the best achievable exponent in the Bayesian probability of error for binary hypothesis testing (see, e.g., [3, Theorem 11.9.1]). Furthermore, if $X_1, X_2, \ldots, X_N$ are i.i.d. random variables, having distribution P with prior probability π_1 and distribution Q with prior probability π_2, the following upper bound holds for the best achievable overall probability of error:

$$P_{\mathrm{e}}^{(N)} \leq \exp\big(-N\,C(P,Q)\big). \tag{13}$$

Definition 6: The *chi-squared divergence* between two probability distributions P and Q, defined on a common set $\mathcal{X}$, is given by

$$\chi^2(P,Q) \triangleq \sum_{x\in\mathcal{X}} \frac{\big(P(x)-Q(x)\big)^2}{Q(x)} = \sum_{x\in\mathcal{X}} \frac{P(x)^2}{Q(x)} - 1\,. \tag{14}$$

The chi-squared divergence is an asymmetric f-divergence where $f(t) = (t-1)^2$ is a convex function with $f(1) = 0$.

For further study of f-divergences and probability metrics, the interested reader is referred to, e.g., [5, Chapter 4], [8, Chapter 2], [12]–[16], [19]–[22], [26]–[30], [33].

II. Improved Bounds for Probability Metrics and f-Divergences

A. Bounds on the Hellinger Distance and Bhattacharyya Parameter

The following proposition introduces a sharpened version of Proposition 1.

Proposition 2: Let P and Q be two probability distributions that are defined on a common set $\mathcal{X}$. Then, the following inequality suggests a tightened version of the inequality in (12)

$$1-\sqrt{1-\big(d_{\mathrm{TV}}(P,Q)\big)^2} \leq \big(d_{\mathrm{H}}(P,Q)\big)^2 \leq \min\left\{1-\exp\left(-\frac{D(P||Q)}{2}\right),\, d_{\mathrm{TV}}(P,Q)\right\} \tag{15}$$

and

$$\max\left\{\exp\left(-\frac{D(P||Q)}{2}\right),\, 1-d_{\mathrm{TV}}(P,Q)\right\} \leq Z(P,Q) \leq \sqrt{1-\big(d_{\mathrm{TV}}(P,Q)\big)^2}. \tag{16}$$

Proof: We start with the proof of the left-hand side of (15). From (4)– (7), and the Cauchy-Schwartz inequality

$$\begin{aligned}
&d_{\mathrm{TV}}(P,Q)\\
&= \frac{1}{2}\sum_{x\in\mathcal{X}} |P(x)-Q(x)|\\
&= \frac{1}{2}\sum_{x\in\mathcal{X}} \Big|\sqrt{P(x)}-\sqrt{Q(x)}\Big|\,\Big(\sqrt{P(x)}+\sqrt{Q(x)}\Big)\\
&\leq \frac{1}{2}\left(\sum_{x\in\mathcal{X}}\Big(\sqrt{P(x)}-\sqrt{Q(x)}\Big)^2\right)^{\frac{1}{2}}\left(\sum_{x\in\mathcal{X}}\Big(\sqrt{P(x)}+\sqrt{Q(x)}\Big)^2\right)^{\frac{1}{2}}\\
&= d_{\mathrm{H}}(P,Q)\cdot\left(1+\sum_{x\in\mathcal{X}}\sqrt{P(x)\,Q(x)}\right)^{\frac{1}{2}}\\
&= d_{\mathrm{H}}(P,Q)\left(2-\big(d_{\mathrm{H}}(P,Q)\big)^2\right)^{\frac{1}{2}}.
\end{aligned} \tag{17}$$

Let $c \triangleq \big(d_{\mathrm{TV}}(P,Q)\big)^2$ and $x \triangleq \big(d_{\mathrm{H}}(P,Q)\big)^2$. By squaring both sides of (17), it follows that $x(2-x)\geq c$, which therefore implies that

$$1-\sqrt{1-c} \leq x \leq 1+\sqrt{1-c}\,. \tag{18}$$

The right-hand side of (18) is satisfied automatically since $0 \leq d_{\mathrm{H}}(P,Q) \leq 1$ implies that $x \leq 1$. The left-hand side of (18) gives the lower bound on the left-hand side of (15). Next, we prove the upper bound on the right-hand side of (15). The use of Jensen's inequality gives

$$\begin{aligned}
&\big(d_{\mathrm{H}}(P,Q)\big)^2\\
&= \frac{1}{2}\sum_{x\in\mathcal{X}}\Big(\sqrt{P(x)}-\sqrt{Q(x)}\Big)^2\\
&= 1-\sum_{x\in\mathcal{X}}\sqrt{P(x)\,Q(x)}\\
&= 1-\sum_{x\in\mathcal{X}} P(x)\sqrt{\frac{Q(x)}{P(x)}}\\
&= 1-\sum_{x\in\mathcal{X}} P(x)\,e^{\frac{1}{2}\log\left(\frac{Q(x)}{P(x)}\right)}\\
&\leq 1-e^{\frac{1}{2}\sum_{x\in\mathcal{X}} P(x)\log\left(\frac{Q(x)}{P(x)}\right)}\\
&= 1-e^{-\frac{1}{2}D(P||Q)}
\end{aligned} \tag{19}$$

and the inequality $\big(d_{\mathrm{H}}(P,Q)\big)^2 \leq d_{\mathrm{TV}}(P,Q)$ is due to [18, Lemma 1]; its (somewhat simplified) proof is as follows:

$$\begin{aligned}\big(d_{\mathrm{H}}(P,Q)\big)^2 &= \frac{1}{2}\sum_{x\in\mathcal{X}}\Big(\sqrt{P(x)}-\sqrt{Q(x)}\Big)^2\\ &= \frac{1}{2}\sum_{x\in\mathcal{X}}|P(x)-Q(x)|\left(\frac{|\sqrt{P(x)}-\sqrt{Q(x)}|}{\sqrt{P(x)}+\sqrt{Q(x)}}\right)\\ &\leq \frac{1}{2}\sum_{x\in\mathcal{X}}|P(x)-Q(x)| = d_{\mathrm{TV}}(P,Q).\end{aligned}$$

The combination of the two upper bounds on the squared Hellinger distance provides the upper bound on the right-hand side of (15). The other bound on the Bhattacharyya parameter in (16) follows from (15) and the simple relation in (7) between the Bhattacharyya parameter and Hellinger distance. ■

Discussion 1: The proof of Proposition 2 is elementary. It is interesting to realize that the sharpened lower bound on the Hellinger distance in terms of the total variation distance, as is given in (15), also follows from the (more involved) lower bound on symmetric f-divergences in (2). To verify this, a combination of (2), (8), (9) gives

$$\begin{aligned}\big(d_{\mathrm{H}}(P,Q)\big)^2 &\geq \big(1-d_{\mathrm{TV}}(P,Q)\big)\cdot\frac{1}{2}\left(1-\sqrt{\frac{1+d_{\mathrm{TV}}(P,Q)}{1-d_{\mathrm{TV}}(P,Q)}}\right)^2\\ &= \frac{1}{2}\Big(\sqrt{1+d_{\mathrm{TV}}(P,Q)}-\sqrt{1-d_{\mathrm{TV}}(P,Q)}\Big)^2\\ &= 1-\sqrt{1-\big(d_{\mathrm{TV}}(P,Q)\big)^2}\end{aligned}$$

which coincides with the left-hand side of the inequality in (15). Similarly, the right-hand side of (16) follows from the equality in (7) and the left-hand side of (15). Hence, it yields that 'half' of Proposition 2 follows from [15, Corollary 5.4], although the proof in this paper is elementary.

Remark 1: Since the total variation distance $d_{\mathrm{TV}}(P,Q)$ and the Hellinger distance $d_{\mathrm{H}}(P,Q)$ are symmetric in P and Q, in contrast to the relative entropy $D(P||Q)$, one can improve the upper bound on the Hellinger distance as follows (see the right-hand side of (15)):

$$d_{\mathrm{H}}(P,Q) \leq \sqrt{\min\left\{1-\exp\left(-\frac{1}{2}\,\min\big\{D(P||Q),\,D(Q||P)\big\}\right),\,d_{\mathrm{TV}}(P,Q)\right\}} \tag{20}$$

and, from (7), the lower bound on the Bhattacharyya parameter on the left-hand side of (16) is improved to

$$Z(P,Q) \geq \max\left\{\exp\left(-\frac{1}{2}\,\min\big\{D(P||Q),\,D(Q||P)\big\}\right),\,1-d_{\mathrm{TV}}(P,Q)\right\}. \tag{21}$$

Remark 2: The bounds in (12) (proved, e.g., in [23]) follow from a loosening of the bounds in (15) by a use of the inequalities $\sqrt{1-x}\leq 1-\frac{x}{2}$ for $x\in[0,1]$, and $e^{-x}\geq 1-x$ for $x\geq 0$.

Remark 3: A comparison of the upper and lower bounds on the Hellinger distance in (15) or the Bhattacharyya parameter in (16) gives the following lower bound on the relative entropy in terms of the total variation distance:

$$D(P||Q) \geq \log\left(\frac{1}{1-\big(d_{\mathrm{TV}}(P,Q)\big)^2}\right). \tag{22}$$

It is noted that (22) also follows from the combination of the last two inequalities in [17, p. 741]. It is tighter than Pinsker's inequality (a.k.a Csiszár-Kemperman-Kullback-Pinsker inequality)

$$D(P||Q) \geq 2\big(d_{\mathrm{TV}}(P,Q)\big)^2$$

when $d_{\mathrm{TV}}(P,Q)\geq 0.893$, having also the advantage of giving the right bound for the relative entropy (∞) when the total variation distance is approached to 1. However, (22) is a slightly looser bound on the relative entropy in comparison to Vajda's lower bound [29] that reads:

$$D(P||Q) \geq \log\left(\frac{1+d_{\mathrm{TV}}(P,Q)}{1-d_{\mathrm{TV}}(P,Q)}\right) - \frac{2d_{\mathrm{TV}}(P,Q)}{1+d_{\mathrm{TV}}(P,Q)}\,. \tag{23}$$

B. A Lower Bound on the Chernoff Information in Terms of the Total Variation Distance

Proposition 3: Let P and Q be two probability distributions that are defined on a common set $\mathcal{X}$. Then, the Chernoff information between P and Q is lower bounded in terms of the total variation distance as follows:

$$C(P,Q) \geq -\frac{1}{2}\,\log\Big(1-\big(d_{\mathrm{TV}}(P,Q)\big)^2\Big). \tag{24}$$

Proof:

$$\begin{aligned} C(P,Q) &\overset{\text{(a)}}{\geq} -\log\left(\sum_{x\in\mathcal{X}}\sqrt{P(x)\,Q(x)}\right) \\ &\overset{\text{(b)}}{=} -\log\, Z(P,Q) \\ &\overset{\text{(c)}}{=} -\log\,\Big(1-\big(d_{\mathrm{H}}(P,Q)\big)^2\Big) \\ &\overset{\text{(d)}}{\geq} -\frac{1}{2}\log\Big(1-\big(d_{\mathrm{TV}}(P,Q)\big)^2\Big) \end{aligned}$$

where inequality (a) follows by selecting the possibly sub-optimal choice $\theta=\frac{1}{2}$ in (10), equality (b) holds by definition (see (6)), equality (c) follows from (7), and inequality (d) follows from the left-hand side of (15). ■

Remark 4: A lower bound on the total variation distance implies a lower bound on the Chernoff information (see Proposition 3), which in turn provides an upper bound on the best achievable Bayesian probability of error for binary hypothesis testing (see, e.g., [3, Theorem 11.9.1] and (13)). For example, lower bounds on the total variation distance in the context of the Poisson approximation were obtained via the use of the Chen-Stein method in [1] and [25]. Another lower bound on the total variation distance appears in [32], followed by Proposition 3 (that was originally introduced in [24, Proposition 5]) to obtain a lower bound on the Chernoff information in the context of the communication problem studied in [32].

C. Bounds on the Chi-Squared Divergence & Related Inequalities for Relative Entropy and Total Variation Distance

Proposition 4: Let P and Q be two probability distributions that are defined on a common set $\mathcal{X}$. Then, the chi-squared divergence between P and Q is lower bounded in terms of the relative entropy as follows:

$$\chi^2(P,Q) \geq e^{D(P||Q)}-1 \tag{25}$$

and, it is also lower bounded in terms of the total variation distance as follows:

$$\chi^2(P,Q) \geq \frac{\big(1+d_{\mathrm{TV}}(P,Q)\big)^{d_{\mathrm{TV}}(P,Q)}}{1-\big(d_{\mathrm{TV}}(P,Q)\big)^2}-1. \tag{26}$$

Furthermore, if $\mathcal{X}$ is a finite set, the following upper bound holds:

$$\chi^2(P,Q) \leq \frac{2\big(d_{\mathrm{TV}}(P,Q)\big)^2}{\min_{x\in\mathcal{X}} Q(x)}\,. \tag{27}$$

Proof: From (14), it follows that

$$\begin{aligned} \chi^2(P,Q) &= \sum_{x\in\mathcal{X}}\frac{P(x)^2}{Q(x)}-1 \\ &= \sum_{x\in\mathcal{X}}\Big\{P(x)e^{\log\left(\frac{P(x)}{Q(x)}\right)}\Big\}-1 \\ &\geq e^{\sum_{x\in\mathcal{X}}P(x)\,\log\left(\frac{P(x)}{Q(x)}\right)}-1 \\ &= e^{D(P||Q)}-1 \end{aligned}$$

where the last inequality follows from Jensen's inequality. This proves the inequality in (25).

The second lower bound on the chi-squared divergence in (26), expressed in terms of the total variation distance, follows from a combination of the first lower bound in (25) with the improvement in [14] of Vajda's inequality:

$$D(P||Q) \geq \log\left(\frac{1}{1-d_{\mathrm{TV}}(P,Q)}\right) - \big(1-d_{\mathrm{TV}}(P,Q)\big)\,\log\big(1+d_{\mathrm{TV}}(P,Q)\big). \tag{28}$$

For the derivation of the upper bound on the chi-squared divergence in (27), note that

$$\begin{aligned}
\chi^2(P,Q) &= \sum_{x\in\mathcal{X}} \frac{\big(P(x)-Q(x)\big)^2}{Q(x)} \\
&\leq \frac{\sum_{x\in\mathcal{X}}\big(P(x)-Q(x)\big)^2}{\min_{x\in\mathcal{X}} Q(x)} \qquad (29)\\
&\leq \frac{\Big(\sum_{x\in\mathcal{X}}|P(x)-Q(x)|\Big)^2}{\min_{x\in\mathcal{X}} Q(x)} \\
&= \frac{4\big(d_{\mathrm{TV}}(P,Q)\big)^2}{\min_{x\in\mathcal{X}} Q(x)} \qquad (30)
\end{aligned}$$

where the last equality follows from (4). However, the upper bound in (27) is twice smaller than (30). In order to prove the tightened upper bound on the chi-squared divergence in (27), we rely on (29), and the following lemma:

Lemma 1: Let

$$d_{\mathrm{loc}}(P,Q) \triangleq ||P-Q||_\infty = \sup_{x\in\mathcal{X}} |P(x)-Q(x)| \tag{31}$$

be the *local distance* between a pair of probability distributions P and Q defined on a set $\mathcal{X}$. Then, the inequality $d_{\mathrm{loc}}(P,Q) \leq d_{\mathrm{TV}}(P,Q)$ holds, which means that the l_∞-norm of $P-Q$ does not exceed *one-half* of its l_1-norm.

Proof: This known inequality follows directly from (3) and (4). ■

As a continuation to the proof of (27), it follows from (29) and Lemma 1 that

$$\begin{aligned}
\chi^2(P,Q) &\leq \frac{\sum_{x\in\mathcal{X}}\big(P(x)-Q(x)\big)^2}{\min_{x\in\mathcal{X}} Q(x)} \\
&\leq \frac{\max_{x\in\mathcal{X}}|P(x)-Q(x)| \cdot \sum_{x\in\mathcal{X}}|P(x)-Q(x)|}{\min_{x\in\mathcal{X}} Q(x)} \\
&\overset{\text{(a)}}{=} \frac{2\,d_{\mathrm{loc}}(P,Q)\,d_{\mathrm{TV}}(P,Q)}{\min_{x\in\mathcal{X}} Q(x)} \\
&\overset{\text{(b)}}{\leq} \frac{2\big(d_{\mathrm{TV}}(P,Q)\big)^2}{\min_{x\in\mathcal{X}} Q(x)}
\end{aligned}$$

where equality (a) follows from (4) and (31) (note that $\mathcal{X}$ is a finite set), and inequality (b) follows from Lemma 1. To conclude, the upper bound on the chi-squared divergence is improved by a factor of 2, as compared to (30), where this improvement is obtained by taking advantage of Lemma 1. ■

Remark 5: Inequality (25) dates back to Dragomir and Gluščević (see [9, Theorem 4]).[1] The lower bound on the chi-squared divergence in (25) significantly improves the Csiszár-Györfi-Talata bound[2] in [6, Lemma 6.3] which states that $\chi^2(P,Q) \geq D(P||Q)$ (note that $e^x \geq 1+x$ for $x \geq 0$).

Remark 6: The transition from (a) to (b) in the derivation of the new upper bound in (27) implies that the improvement by a factor of 2 that is obtained there, as compared to (30), can be further enhanced under a mild condition. Specifically, a further improvement is obtained if the ratio $\frac{d_{\mathrm{loc}}(P,Q)}{d_{\mathrm{TV}}(P,Q)}$, which according to Lemma 1 is no more than 1, is strictly below 1 (for such possible examples, the reader is referred to [26, Section 4]); in this case, the improvement over the upper bound on the chi-squared divergence in (30) is by a factor of $\frac{2\,d_{\mathrm{TV}}(P,Q)}{d_{\mathrm{loc}}(P,Q)}$.

[1] Inequality (25) is missing a proof in [9]; it was recently proved in [27, Theorem 3.1], and it was derived independently in this work (before being aware of [9] and [27]).

[2] As a historical note, Györfi was acknowledged for pointing out the inequality $\chi^2(P,Q) \geq D(P||Q)$ in [6, Lemma 6.3]; this inequality was earlier stated in [4, Lemma 4] under a redundant requirement (see also [7, Lemma A.7], stated with a variant of this requirement).

The following is a sort of a reverse of Pinsker's inequality:

Corollary 1: Let P and Q be two probability distributions that are defined on a common finite set $\mathcal{X}$. Then, the following inequality holds:

$$D(P||Q) \leq \log\left(1 + \frac{2\big(d_{\mathrm{TV}}(P,Q)\big)^2}{\min_{x\in\mathcal{X}} Q(x)}\right). \tag{32}$$

Proof: This result follows from the bounds on the chi-squared divergence in (25) and (27). ■

Remark 7: The bound in (32) improves the bound that follows by combining Csiszár-Györfi-Talata bound in [6, Lemma 6.3] (see Remark 5) and the bound in (30). This combination gives the Csiszár-Györfi-Talata bound

$$\min_{x\in\mathcal{X}} Q(x)\, D(P||Q) \leq 4\big(d_{\mathrm{TV}}(P,Q)\big)^2. \tag{33}$$

The improvement that is suggested in (32) over (33) is twofold: the logarithm on the right-hand side of (32) follows from the lower bound on the chi-squared divergence in (25) (as compared to the inequality $\chi^2(P,Q) \geq D(P||Q)$ in [6, Lemma 6.3]); another improvement, obtained by a replacement of the factor 4 on the right-hand side of (33) by a factor 2 inside the logarithm on the right-hand side of (32), follows from the improvement of the upper bound on the chi-squared divergence in (27) over the bound in (30).

Note that when the distributions P and Q are close enough in total variation, the upper bounds on the relative entropy in (32) and (33) scale like the square of the total variation distance (although the former bound improves the latter bound by a factor of 2).

Remark 8: The following inequality has been recently introduced by Verdú: [30]:

$$d_{\mathrm{TV}}(P,Q) \geq \left(\frac{1-\beta}{\log\frac{1}{\beta}}\right) D(P||Q) \tag{34}$$

where $\beta^{-1} \triangleq \sup_{x\in\mathcal{X}} \frac{\mathrm{d}P}{\mathrm{d}Q}(x)$. The reader is also referred to [11, Lemma 3.10] where a related inequality is provided.

Remark 9: The lower bound on the chi-squared divergence in (26) is looser than the bound in (25) (due to the additional use of the inequality in (28) for the derivation of (26)); nevertheless, the bound in (26) is expressed in terms of the total variation distance, whereas the bound in (25) which is expressed in terms of the relative entropy.

Remark 10: As an addition to Proposition 4, a parameterized upper bound on the chi-squared divergence is introduced in [15, Corollary 5.6] where this bound is expressed in terms of some power divergences.

Remark 11: A related problem to the result in Corollary 1 has been recently studied in [2]. The problem studied in [2] is the characterization of $D^*(\varepsilon, Q)$ for an arbitrary distribution Q, defined to be the infimum of $D(P||Q)$ over all distributions P that are at least ε-far away from Q in total variation. It is demonstrated in [2] that $D^*(\varepsilon, Q)$ scales like $C\varepsilon^2 + O(\varepsilon^3)$ for a certain constant C (with explicit upper and lower bounds on C). For the case where P and Q are defined on a common finite set $\mathcal{X}$, the scaling in ε^2 (for $\varepsilon \ll 1$) is supported by the combination of Corollary 1 and Pinsker's inequality. Corollary 1 further implies that (even *not necessarily* in the limit where $\varepsilon \to 0$)

$$D^*(\varepsilon, Q) \triangleq \inf_{P:\, d_{\mathrm{TV}}(P,Q)\geq\varepsilon} D(P||Q) \leq \log\left(1 + \frac{2\varepsilon^2}{\min_{x\in\mathcal{X}} Q(x)}\right).$$

D. Bounds on the Capacitory Discrimination in Terms of the Total Variation Distance

The capacitory discrimination, introduced by Topsøe [28] and further studied in [10] and [15], is a probability metric which forms an f-divergence. It is defined as follows:

Definition 7: Let P and Q be two probability distributions that are defined on a common set $\mathcal{X}$. The capacitory discrimination is given by

$$\overline{C}(P,Q) \triangleq D\left(P \,||\, \frac{P+Q}{2}\right) + D\left(Q \,||\, \frac{P+Q}{2}\right). \tag{35}$$

Due to the parallelogram identity for relative entropy, it follows that

$$\overline{C}(P,Q) = \min_{R} \big\{ D(P||R) + D(Q||R) \big\}$$

where the minimization is taken over all the probability distributions R.

Proposition 5: The capacitory discrimination is lower bounded in terms of the total variation distance as follows:

$$\overline{C}(P,Q) \geq 2\, D\left(\frac{1 - d_{\mathrm{TV}}(P,Q)}{2} \,||\, \frac{1}{2} \right) \tag{36}$$

where $D(p||q) \triangleq p \log\left(\frac{p}{q}\right) + (1-p)\log\left(\frac{1-p}{1-q}\right)$ for $p, q \in [0,1]$ (with the convention that $0 \log 0 = 0$). Furthermore, if $\mathcal{X}$ is a finite set, then it satisfies the following upper bound in terms of the total variation distance:

$$\overline{C}(P,Q) \leq 2 \, \log\left(1 + \frac{\big(d_{\mathrm{TV}}(P,Q)\big)^2}{\min_{x \in \mathcal{X}} \big(P(x) + Q(x)\big)} \right). \tag{37}$$

Proof: In [15, p. 119], the capacitory discrimination is expressed as an f-divergence where the convex function f is given by $f(x) = x \log x - (x+1)\log(1+x) + 2\log 2$, for $x > 0$, with $f(1) = 0$. Although the capacitory discrimination in (35) is symmetric with respect to P and Q, the above function f is *asymmetric* since $f(x) \neq x f\big(\frac{1}{x}\big)$ (see Definition 2). In order to apply the lower bound in (2) (see [13] and [15, Corollary 5.4]), we first need to find a *symmetric* convex function f with $f(1) = 0$ which satisfies the equality $\overline{C}(P,Q) = D_f(P||Q)$. It can be verified that the proper symmetric function f is given by

$$f(x) = x \log x - (x+1)\log(1+x) + (x+1)\log 2, \quad x > 0. \tag{38}$$

Consequently, the combination of (2) and (38) implies that

$$\begin{aligned}
\overline{C}(P,Q) &\geq \big(1 - d_{\mathrm{TV}}(P,Q)\big)\, f\left(\frac{1 + d_{\mathrm{TV}}(P,Q)}{1 - d_{\mathrm{TV}}(P,Q)} \right) \\
&= \big(1 + d_{\mathrm{TV}}(P,Q)\big)\, \log\big(1 + d_{\mathrm{TV}}(P,Q)\big) + \big(1 - d_{\mathrm{TV}}(P,Q)\big)\, \log\big(1 - d_{\mathrm{TV}}(P,Q)\big) \\
&= 2 \left[\log 2 - h\left(\frac{1 - d_{\mathrm{TV}}(P,Q)}{2} \right) \right] \\
&= 2\, D\left(\frac{1 - d_{\mathrm{TV}}(P,Q)}{2} \,||\, \frac{1}{2} \right).
\end{aligned}$$

The last equality holds since $D(p||\frac{1}{2}) = \log 2 - h(p)$ for $p \in [0,1]$ where h denotes the binary entropy function. This proves the lower bound in (36). The derivation of the upper bound in (37) relies on a combination of (32) (see Corollary 1), and the equality

$$d_{\mathrm{TV}}\left(P, \frac{P+Q}{2} \right) = d_{\mathrm{TV}}\left(Q, \frac{P+Q}{2} \right) = \frac{d_{\mathrm{TV}}(P,Q)}{2}.$$

■

Discussion 2: The lower bound on the capacitory discrimination in (36), expressed in terms of the total variation distance, forms a closed-form expression of the bound by Topsøe in [28, Theorem 5]. This bound is given by

$$\overline{C}(P,Q) \geq \sum_{\nu=1}^{\infty} \frac{\big(d_{\mathrm{TV}}(P,Q)\big)^{2\nu}}{\nu(2\nu-1)}. \tag{39}$$

The equivalence of (36) and (39) follows from the power series expansion of the binary entropy function (to the natural base) around one-half (see [31, Lemma 2.1]):

$$h(x) = \log 2 - \sum_{\nu=1}^{\infty} \frac{(1-2x)^{2\nu}}{2\nu(2\nu-1)}, \quad \forall x \in [0,1]$$

which yields that

$$\begin{aligned}
&\sum_{\nu=1}^{\infty} \frac{\left(d_{\mathrm{TV}}(P,Q)\right)^{2\nu}}{\nu(2\nu-1)} \\
&= 2\left[\log 2 - h\left(\frac{1-d_{\mathrm{TV}}(P,Q)}{2}\right)\right] \\
&= 2D\left(\frac{1-d_{\mathrm{TV}}(P,Q)}{2}\,||\,\frac{1}{2}\right).
\end{aligned}$$

Note, however, that the proof here is much shorter than the proof of [28, Theorem 5] (which relies on properties of the triangular discrimination in [28] and previous theorems of this paper), and it also leads directly to a closed-form expression of this bound. Consequently, one concludes that the lower bound in [28, Theorem 5] is a special case of (2) (see [13] and [15, Corollary 5.4]), which provides a lower bound on an arbitrary symmetric f-divergence in terms of the total variation distance.

The upper bound on the capacitory discrimination in (37) is new, and it is based on Corollary 1 which provides an improvement of the Csiszár-Györfi-Talata bound (see Remarks 5 and 7).

III. Summary

Derivation of tight bounds for probability metrics and f-divergences is considered in this paper. In some cases, existing recent bounds are reproduced by elementary or simplified proofs, and in some other cases, elementary proofs provide significant improvements. The contributions made in this work are outlined in the following:

- Upper and lower bounds on both the Hellinger distance and the Bhattacharyya parameter are expressed in terms of the total variation distance and relative entropy (see Proposition 2). The lower bound on the Hellinger distance and the upper bound on the Bhattacharyya parameter are not new; nevertheless, their proofs here are simple and elementary (see Discussion 1). The other two bounds are new.
- A new lower bound on the Chernoff information is expressed in terms of the total variation distance (see Proposition 3). It has been recently applied in [32] with a reference to the un-published bound in [24, Proposition 5].
- Three bounds on the chi-squared divergence are introduced in Proposition 4. The first lower bound in (25) dates back to Dragomir and Gluščević [9] (see Remark 5). A second lower bound on the chi-squared divergence is derived in terms of the total variation distance (see (26)). The upper bound on the chi-squared divergence in (27) is new as well, and it suggests an improvement over the bound in (30) by a factor of 2 (according to Remark 6, this gain can be further improved under a mild condition).
- The improvements of the bounds on the chi-squared divergence (as outlined in the previous item) lead to a new improved upper bound on the relative entropy in terms of the total variation distance when the two probability distributions are defined on a common finite set (see Corollary 1, followed by Remarks 7–11). This forms a kind of a reverse of Pinsker's inequality where the two distributions are defined on a finite set, and it improves the Csiszár-Györfi-Talata bound in (33).
- Bounds on the capacitory discrimination are provided in terms of the total variation distance (see Proposition 5). The lower bound on the capacitory discrimination forms a closed-form expression of the bound by Topsøe in [28, Theorem 5]; its proof, however, is more simple, and it does not involve properties of the triangular discrimination that are used in its original proof in [28]. The upper bound on the capacitory discrimination in (37) is new (see Discussion 2, addressing the bounds in this item).

Historical Background

The material presented in this paper partially appears in the un-published manuscript [24] (parts of [24] have later been published in [25] and [26], without any overlap with the bounds introduced in this paper). A recent progress in this work stimulated the writing of the present paper where the un-published results in [24, Propositions 4, 5] served as a starting point.

Acknowledgment

This research work was supported by the Israeli Science Foundation (ISF), grant number 12/12.

REFERENCES

[1] A. D. Barbour, L. Holst and S. Janson, *Poisson Approximation*, Oxford University Press, 1992.

[2] D. Berend, P. Harremoës and A. Kontorovich, "Minimum KL-divergence on complements of L_1 balls," to appear in the *IEEE Trans. on Information Theory*, vol. 60, 2014.

[3] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, John Wiley and Sons, second edition, 2006.

[4] I. Csiszár, "Large-scale typicality of Markov sample paths and consistency of MDL order estimators," *IEEE Trans. on Information Theory*, vol. 48, no. 6, pp. 1616–1628, June 2002.

[5] I. Csiszár and P. C. Shields, *Information Theory and Statistics: A Tutorial*, Foundations and Trends in Communications and Information Theory, vol. 1, no. 4, pp. 417–528, 2004.

[6] I. Csiszár and Z. Talata, "Context tree estimation for not necessarily finite memory processes, via BIC and MDL," *IEEE Trans. on Information Theory*, vol. 52, no. 3, pp. 1007–1016, March 2006.

[7] I. Csiszár and Z. Talata, "Consisitent estimation of the basic neighborhood of Markov random fields," *Annals od Statistics*, vol. 34, no. 1, pp. 123–145, May 2006.

[8] A. DasGupta, *Asymptotic Theory of Statistics and Probability*, Springer Texts in Statistics, 2008.

[9] S. S. Dragomir and V. Gluščević, "Some inequalities for the Kullback-Leibler and χ^2-distances in information theory and applications," *Tamsui Oxford Journal of Mathematical Sciences*, vol. 17, no. 2, pp. 97–111, 2001.

[10] D. M. Endres and J. E. Schindelin, "A new metric for probability distributions," *IEEE Trans. on Information Theory*, vol. 49, no. 7, pp. 1858–1860, July 2003.

[11] E. Even-Dar, S. M. Kakade and Y. Mansour, "The value of observation for monitoring dynamical systems," *Proceedings of the International Joint Conference on Artificial Intelligence*, pp. 2474–2479, Hyderabad, India, January 2007.

[12] A. L. Gibbs and F. E. Su, "On choosing and bounding probability metrics," *International Statistical Review*, vol. 70, no. 3, pp. 419–435, 2002.

[13] G. L. Gilardoni, "On the minimum f-divergence for given total variation," *Comptes Rendus Mathematique*, vol. 343, no. 11–12, pp. 763–766, 2006.

[14] G. L. Gilardoni, "On Pinsker's and Vajda's type inequalities for Csiszár's f-divergences," *IEEE Trans. on Information Theory*, vol. 56, no. 11, pp. 5377–5386, November 2010.

[15] A. Guntuboyina, S. Saha, and G. Schiebinger, "Sharp inequalities for f-divergences," *IEEE Trans. on Information Theory*, vol. 60, no. 1, pp. 104–121, January 2014.

[16] P. Harremoës, "Information topologies with applications," *Entropy, Search, Complexity*, vol. 16, pp. 113–150, 2007.

[17] W. Hoeffding and J. Wolfowitz, "Distinguishability of sets of distributions," *Annals of Mathematical Statistics*, vol. 29, no. 3, pp. 700–718, September 1958.

[18] C. Kraft, "Some conditions for consistency and uniform consistency of statistical procedures," *University of California Publications in Statistics*, vol. 1, pp. 125–142, 1955.

[19] F. Liese and I. Vajda, "On divergences and informations in statistics and information theory," *IEEE Trans. on Information Theory*, vol. 52, no. 10, pp. 4394–4412, October 2006.

[20] H. Mostafaei and S. Kordnourie, "Probability metrics and their applications," *Applied Mathematical Sciences*, vol. 5, no. 4, pp. 181–192, 2011.

[21] F. Nielsen, "An information-geometric characterization of Chernoff information," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 269–272, March 2013.

[22] V. V. Prelov and E. C. van der Meulen, "Mutual information, variation, and Fano's inequality," *Problems of Information Transmission*, vol. 44, no. 3, pp. 185–197, September 2008.

[23] R. D. Reiss, *Approximate Distributions of Order Statistics with Applications to Non-Parametric Statistics*, Springer Series in Statistics, Springer-Verlag, 1989.

[24] I. Sason, "An information-theoretic perspective of the Poisson approximation via the Chen-Stein method," un-published manuscript, *arXiv:1206.6811v4*, 2012.

[25] I. Sason, "Improved lower bounds on the total variation distance for the Poisson approximation," *Statistics and Probability Letters*, vol. 83, no. 10, pp. 2422–2431, October 2013.

[26] I. Sason, "Entropy bounds for discrete random variables via maximal coupling," *IEEE Trans. on Information Theory*, vol. 59, no. 11, pp. 7118–7131, November 2013.

[27] A. Sayyareh, "A new upper bound for Kullback-Leibler divergence," *Applied Mathematical Sciences*, vol. 5, no. 67, pp. 3303–3317, 2011.

[28] F. Topsøe, "Some inequalities for information divergence and related measures of discrimination," *IEEE Trans. on Information Theory*, vol. 46, pp. 1602–1609, July 2000.

[29] I. Vajda, "Note on discrimination information and variation," *IEEE Trans. on Information Theory*, vol. 16, no. 6, pp. 771–773, November 1970.

[30] S. Verdú, "Total variation distance and the distribution of the relative information," presented at the *9th Workshop on Information Theory & Applications (ITA 2014)*, La Jolla, San-Diego, California, USA, February 9–14, 2014.

[31] G. Wiechman and I. Sason, "Parity-check density versus performance of binary linear block codes: new bounds and applications," *IEEE Trans. on Information Theory*, vol. 53, no. 2, pp. 550–579, February 2007.

[32] A. D. Yardi. A. Kumar, and S. Vijayakumaran, "Channel-code detection by a third-party receiver via the likelihood ratio test," accepted to the *2014 IEEE International Symposium on Information Theory*, Honolulu, Hawaii, USA, July 2014.

[33] Z. Zhang, "Estimating mutual information via Kolmogorov distance," *IEEE Trans. on Information Theory*, vol. 53, no. 9, pp. 3280–3282, September 2007.